arXiv:1701.07657v1 [cs.AI] 26 Jan 2017

Logic Programming Petri Nets

Giovanni Sileno

[1] University of Amsterdam, Leibniz Center for Law, the Netherlands
[2] Télécom ParisTech, TII/TSI, Paris, France
g.sileno@uva.nl

Abstract. With the purpose of modeling, specifying and reasoning in an integrated fashion with procedural and declarative aspects (both commonly present in cases or scenarios), the paper introduces Logic Programming Petri Nets (LPPN), an extension to the Petri Net notation providing an interface to logic programming constructs. Two semantics are presented. First, a hybrid operational semantics that separates the process component, treated with Petri nets, from the constraint/terminological component, treated with Answer Set Programming (ASP). Second, a denotational semantics maps the notation to ASP fully, via Event Calculus. These two alternative specifications enable a preliminary evaluation in terms of reasoning efficiency.

Keywords: Logic programming, Declarative Rules, Reactive Rules, Reasoning, Simulation, Terminology, Ontology, Answer Set Programming

Introduction

Despite their widespread use in domains as computer science, electronics, business process modeling and biology, Petri nets are generally considered not to be enough expressive for reasoning purposes; in effect, they do not refer explicitly to any informational/representational concept. In their simplest form, tokens are indistinct, and do not transport any data. Nevertheless, usually modelers introduce labels to set up a correspondence between the *modeling* entities and the *modeled* entities. This practice enables them to read the results of the simulation in reference to the modeled system, and therefore it becomes *functional to the use* of the notation, although it is not a requirement for the execution in itself.

Further interaction is possible if these labels are processed according an additional formalism, i.e. if tokens are labeled with some sort of declarative language, adequately integrated with the standard operational semantics. For its generality, applicative value and popularity, the most important notation following this approach is certainly *Coloured Petri Nets* (CPNs) [14], which, for many aspects, is a descendant of *Predicate/Transition Nets* [13]. Other examples exist, however, tailored to specific purposes; for instance, *Simple Logic Petri Nets* (SLPNs), used to perform model-checking of agent programs [5] and *Petri Nets with Data* (DPN), used to perform conformance checking on deviations of data produced by processes [21]. The present paper introduces yet an additional notation: *Logic*

Programming Petri Nets (LPPNs), with the purpose of investigating an alternative solution to interrelated long-standing issues in the knowledge representation domain: to satisfy the requirement of working with causal models and terminological knowledge; to confront the general confusion about declarative and reactive rules. The paper proceeds as follows. § 1 introduces the general problems, the motivation underlying this proposal, and arguments in support to this new notation (§ 1.4). § 2 presents the notation informally, while its formalization and semantics are presented in § 3. First results of evaluation are reported in § 4. A note on further developments ends the paper.

1 Representing and processing knowledge

In psychology, knowledge is traditionally divided between *procedural knowledge* (the *know-how*), often implicit and concerned by how to perform actions, and *declarative knowledge* (the *know-what*), which is conscious and can be communicated in a verbal form [1]. In programming, these two aspects can be fairly associated with two orthogonal paradigms: *imperative* and *declarative*. Following the imperative paradigm, we "command" the computer via a list of *instructions*, i.e. a sequence of actions to be performed in order to obtain a certain goal (which remains hidden to the machine). In contrast, following the declarative paradigm, here exemplified by *declarative logic programming*, we provide the computer with a certain *knowledge*, consisting of *facts* and *rules*, and rather then executing it, the machine is required to *solve* two types of problems with it. The first type of problem involves querying the knowledge, i.e. asking the solver to find one or all elements which satisfy an explicit *goal* expression. This processing is usually implemented in *backward chaining*: the solver attempts to establish whether there is a *proof* that entails the goal from the known facts. The second type of problem involves providing all the facts that are entailed by the knowledge, asking the solver to produce all the possible consequences from the known facts, following and constrained by the rules given as input. This is usually associated with *forward chaining*. Languages whose solvers can, in practice, be seen as associated with these types of problems/processing are respectively Prolog and *answer set programming* (ASP).[3]

1.1 Procedural knowledge is about *transients*

Reinterpreting the imperative paradigm from a knowledge point of view, we could say that, from a modeler's perspective, procedural knowledge is arguably

[3] This is not necessarily correct from an implementation point of view. Both SLD/SLDNF resolution (Prolog) and DPLL (ASP) are based on backward chaining. In DPLL, however, all variables are grounded, and all intermediate atoms generated in the search are collected in *stable models*; without defining any goal, all the worlds that are implied by the input knowledge are returned as output. From an external perspective, this is the same result we would associate with forward chaining. The intuition that there is a relation between ASP and forward chaining is confirmed in ASPeRiX [20].

a natural way to specify *causal mechanisms*: with a correct program, if the machine follows the instructions, it causes the intended result.[4] On closer inspection, however, we observe that it specifies the sequential aspect of the process, rather than essential causal dependencies.[5] A more correct interpretation would be to say that procedural knowledge describes a *transient*[6], i.e. an intermediary phase which goes from an input to an output state. Note that in non-deterministic settings, a stable dynamic system may exhibit multiple outcome-equivalent transients. Returning to the previous terminology, we could then say that, in procedural characterizations, temporal contingent aspects are added on top of structural, causal aspects, which remain tacit.

1.2 Declarative knowledge is about *steady states*

Declarative knowledge is unanimously considered to be a powerful tool for modeling definitions and constraints holding at the terminological and ontological level. Continuing the metaphor with physics, we could say that this type of representations focuses on *steady-state* aspects of reality: how entities are related, *after* the transient has passed. This view is consistent with one of the requirements associated with pure declarative programming: the absence of (or better, the transparency to) *side-effects*. The world that is represented by this knowledge has to be a still world.

Representing change But what about the representation of change? If you want to represent a transition in a declarative way, you can consider *snapshots* of the arrangements holding before and after the transition—possibly labeled with a sort of timestamp. This is, very simplistically, the principle behind *situation calculus* [23,29], *event calculus* [17,30], and *fluent calculus* [32]. Using appropriate axioms, you can create and reason about the relations between these snapshots in a way such that they correspond to the natural relation between the moments they refer to. Unfortunately, these methods do not provide the most natural visualization to the modeler.

Rather than trying to project one dimension on the other, an alternative tradition in AI and logic proposes to consider causality as a primitive notion, not to be reduced to simpler concepts. This approach is, for instance, behind

[4] Contrary to the psychological case, in the computational domain, procedural knowledge seems to be completely explicit, but in reality, it overlooks what is done by the machine *below* the instruction level. Primitive actions can be seen in practice as declarative goals given to the machine to be 'solved'.

[5] Consider the recipe to cook a dish of spaghetti: for the final outcome, it really does not matter if you put the salt in the water before the water starts boiling or after. Vice versa, it *does* matter for the final outcome whether you add it while you are cooking the pasta or after you drain it.

[6] In acoustics and in electronics, the *transient* is a short-duration, high-frequency oscillation occurring when there is great change in amplitude. Non-digital systems never pass abruptly from one state the other, but follow a smoother, continuous pattern, up to eventually reach a new equilibrium state.

the idea of all *Action languages* [11]. Even when the dichotomy is made clear, however, operationalizations of these languages often result in compiling action programs to logic programs [10,9], returning to 'snapshot-handling' solutions again. Is this the best computational solution available?

Intuitively, leaving process analysis to procedural descriptions should be a better choice: as procedural knowledge exploits causal computational mechanisms, it does not only *re-present*, but also *re-create* the process object, transforming the question from what the referent *should* be (characteristic of logic), to what it *is* (characteristic of simulation, intended as model-execution). Recent research adopting a similar approach shows that there is a great computational advantage at stake (see e.g. [26]).

1.3 Logical conditionals and causal dependencies

As convincingly noted by [16], there is a widespread confusion in cognitive science and computational disciplines around the notion of *rules*, primarily because of the lack of neat distinction between *declarative rules* and *reactive rules*. A prototypical example of this problem is the use of the term "production rule". In (formal) linguistics, production rules define the grammar of a language, i.e. the combinatorial system defining valid linguistic elements. In the tradition of expert systems and active databases, production rules are reactive rules: rules guiding actions or interventions, in response to certain conditions. These interpretations deal with two perspectives: one concerning *objects* (e.g. symbols) and their relations, and the other concerning patterns of actions, mechanisms, or *processes* (involving objects)—the declarative and procedural dichotomy again. Formal logic is the prototypical domain of the first perspective, just as process modeling is of the second.

A possible reason for this confusion could be the tension between notions as *logic conditional* and *causal dependency*. While *logic* is naturally biased towards the first operator, as it is established, well-known, and tailored for abstract worlds, and *engineering* favours the second notion because it is crucial for practical purposes, other domains have a less neat positioning. For instance, *law* concerns a system of norms (which, in abstract, could be approached with logic), contextualized and applied to a social system (i.e. dealing with a continuous flow of events). This encounter is made explicit in *cases*.[7] Cases, however, are crucial not only in legal activities, but in practically all human activities. From a discipline perspective, the applicative province of a notation that considers both causal and logic primitives touches upon the *requirement engineering* domain, as it naturally deals with *use cases*, *scenarios*, etc.

[7] The LPPN notation has been first introduced in [31], as a solution to deal with the alignemnt of representations of law (norms), representations of implementations of law (services, i.e. business processes) and representation of action (behavioural patterns, possibly intentionally charaterized).

1.4 Why Logic Programming Petri Nets?

A proper treatment of a case or scenario is based on two requirements: on the one hand, to capture the relationships expressed via language components—i.e. the symbols that we use to *represent* the target system, interrelated at the formal level in a local ontology, pertinent to the scenario; on the other hand, to *reproduce*, via the execution of the model, the same changes exhibited by the target system. The first aspect is the typical application domain of *knowledge bases*, and, more recently, of *semantic ontologies*, introduced with the explicit purpose of defining reusable concepts and relationships ('terms') to describe and represent specific areas of concern. The second aspect is the main concern of *process modeling*, used both for descriptive and prescriptive purposes. Unfortunately, methodologies associated with one of the two aspects generally have a limited treatment of the other component, or they refer to specific mediating machinery to deal with; therefore, a proper treatment of scenarios is not a direct endeavor.

The fundamental dichotomy between logic conditionals and causal dependencies motivates the introduction of a new notation. The idea, in short, is to explicitly distinguish the declarative and procedural components in the model, so as to allow us to treat them separately, recomposing the corresponding results at a higher level. From a conceptual point of view, this choice moves towards mirroring the common-sense distinction between *objects* and *events* [7], roughly reflecting the prototypical use of the *noun/verb* categories in language (cf. cognitive linguistics studies, e.g. [15]). Operationally, this means that the responsibility of the logic dimension may go to a logic solver, and the responsibility of the process dimension to a process model executor.

Because of their simplicity, their visual power, and their wide-spread use, we have decided to consider Petri Nets for the process dimension, while for the logic dimension, we turned to *Answer Set Programming* (ASP), a relatively recent logic programming language, to operationalize a first proof-of-concept of the approach; other solutions are, in principle, possible (e.g. *semantic web* languages), but are outside the scopes of this paper.

2 Informal presentation to LPPNs

Many extensions and variations to the basic Petri Net notation have been proposed in the literature. Some of these—e.g. *Coloured Petri Nets* (CPN) [14] or *Simple Logic Petri Nets* (SLPN) [5]—propose to integrate the descriptive function of labels with the executable nature of the net. The basic idea is that, in order to enable a full computational operationalization, labels should be written in some sort of declarative language, fully integrated in the operational semantics. We follow the same idea, although we start from different requirements.[8]

[8] CPNs are very powerful in terms of expressiveness, and they are applied in many domains; however, they introduce many details which are unimportant in our setting (e.g. expressions on arcs), and this overload does not pay off for its advantages.

If a language is required to model a world and its mechanisms, there should be a correspondence between the *modeling* entities and the *modeled* entities belonging to the reference domain, and this correspondence should be maintained both at the level of entities and mechanisms (ontology as existence) and at the level of concepts and terminology (ontology as shared conceptualization). These are basically the requirements that motivates the introduction of *Logic Programming Petri Nets* (LPPNs). In short, LPPNs extend basic Petri Nets—allowing the specification of procedural mechanisms—with (a) a Prolog-like expression notation for labels, attached on places and transitions (and not on arcs, as e.g. in SLPNs) and (b) with declarative net components—allowing the specification of purely declarative mechanisms.[9]

2.1 Net components

LPPNs reuse the standard components of Petri Nets in a specific way:

- *tokens* are entities that currently exist, described with propositional labels;
- *places*, containers for tokens, function as a sort of *local relational databases*;
- *transitions*, consuming and producing tokens when fired, function as a sort of mechanisms of *retraction* and *assertion*.

Additionally, LPPNs include a new type of component:

- *logic operator nodes*, used to specify logic constraints or for logic compositions binding places, or transitions.

2.2 Labeling

Two levels of specification are available for the labels: the *token-instance* level (for tokens), and the *type-class* level (for places and transitions)

- the *label of a token* defines the propositional content of the associated proposition (a *fact* of a *knowledge base*);
- the *label of a place* defines the relational model of the local database associated with that place;
- the *label of a transition* specifies the parameters of the *transition events* that may be fired.

SLPNs have been introduced for a different purpose, as intermediate notation to transform an AgentSpeak(L) script into an ASP program for model-checking. Beyond these considerations, both notations neglect the declarative (e.g. terminological) dimension of the modeled system.

[9] A similar line of research can be found in the *database systems* literature, with the distinction between *deductive rules* and *active rules*. An important contribution of this track is *transaction logic programming* (TLP) [6]. As the LPPN notation takes *locality* into account, and provides a visual notation, it can be thought as providing a higher abstraction with respect to these proposals. However, further investigation is needed to fully appreciate their correspondences.

2.3 Procedural mechanisms

An example of a sub-net built upon a procedural mechanism is illustrated in Fig. 1. In (a), two facts hold: `p1(a1)` and `p2(a2, b1)`. The transition is not enabled to fire because the *variable binding* required by that event (same parameter for `A`) is not satisfied by the present inputs. In (b), an additional fact `p2(a1, b2)` is present, enabling the transition, which fires. Because of the firing, in (c), a new fact is *forged* with the consumed elements, following the label of the output place.

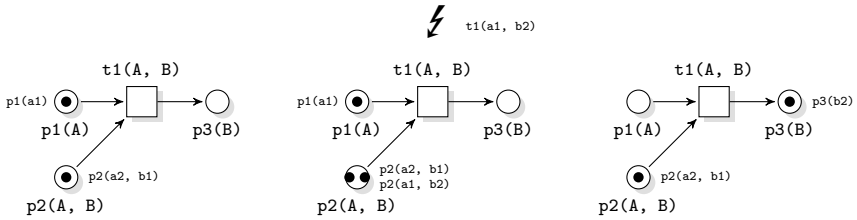

(a) not enabled transition, before firing

(b) enabled transition and firing

(c) the previously enabled transition has fired

Fig. 1: Example of a LPPN *procedural* component and its execution.

2.4 Declarative mechanisms

The LPPN notation adds logic operator nodes, which apply on places or on transitions. An example of a sub-net with logic operator nodes working on places (small black squares) is given in Fig. 2. These are used to create logic compositions of places (via operators as `NEG`, `AND`, `OR`, etc). or to specify logic interdependencies (via the logic conditional `IMPLIES`).

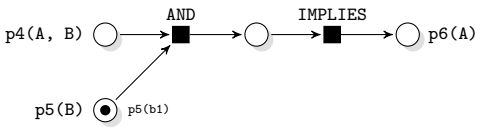

Fig. 2: Example of a LPPN with a *declarative* component defined on *places*, corresponding to the Prolog/ASP code: `p6(A) :- p4(A, B), p5(B). p5(b1).`

Similarly, transitions may be connected declaratively via logic operators working on transitions (black circles) as in Fig. 3. These connections may be interpreted as channels enabling *instantaneous propagation* of firing. In this case, we do not have the usual logic operators, just implication or possibly bi-implication, because, for the interleaving semantics, only one source transition may fire per step. The sense of the arrow is, in principle, sufficient to specify the direction of the relation.

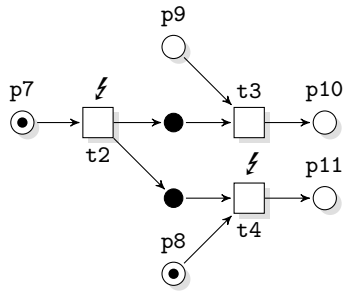

Fig. 3: Example of a LPPN with a *declarative* component defined on *transitions*, instantaneously propagating the firing where possible (in this case, `t2` and `t4` fire together).

Operationally, these declarative components will be treated integrating the semantics of *answer set programming* (ASP) [22]. This was a natural choice because process execution exhibits a prototypical 'forward' nature, and ASP can be interpreted as providing forward chaining (cf. § 1).

Answer Set Programming ASP is a declarative programming paradigm [22] based on the *stable-model semantic* [12], oriented towards difficult (NP-hard) search problems.[10] It is used to model and solve problems belonging to a wide range of applications, and has reached production level maturity. In ASP, similarly to Prolog, the programmer models a problem in terms of rules and facts, instead of specifying an algorithm. The resulting code is given as input to a solver, which returns multiple *answer sets* or *stable models* which satisfy the problem. The main operational difference to Prolog is that all variables are grounded before performing search, and unlike SLDNF resolution, ASP solver algorithms always terminate. Additionally, it provides distinct primitive operators for *default negation* ("I don't know whether p is the case") and *strong negation* ("I know that p is not the case"), which is crucial in all non-monotonic knowledge contexts (i.e. where practical knowledge is paramount: common-sense reasoning, law, medicine, etc.).

3 Formalization

This section provides a formal presentation of *Logic Programming Petri Nets* (LPPN) notation. Despite the different notation, we start similarly to the formalization of *Simple Logic Petri Nets* (SLPN) [5]. However, when addressing the denotational semantics, instead of considering *ad-hoc* transformations to ASP like these authors do (cf. also [3,2]), we explicitly refer to *event calculus* (EC) [17,30], prefiguring future integration with existing literature on linguistic semantics (e.g. [18]).

[10] Stable-model semantics apply ideas of *auto-epistemic* logic [24] and *default logic* [28].

3.1 Labeling language

Definition 1 (Term). *Let CONST be a finite set of* constants*, VAR a finite set of* variables*, and FUNC a finite set of* function symbols*. The (possibly infinite) set of* terms *TERMS consists of:*

- *all constants CONST,*
- *all variables VAR,*
- *all associations of function symbols FUNC to terms, in the form* $f(t_1, \ldots, t_n)$*, with* $f \in FUNC$ *and* $t_1, \ldots, t_n \in TERMS$*.*

Following logic programming languages, we will use strings in lower-case for constants (e.g. `john`) and strings starting with an upper-case letter for variables (e.g. `Person`).

Definition 2 (Ground Term). *A* ground term *is a term not containing variables.*

Definition 3 (Atom). *Given a finite set of* predicate symbols *PRED, an atom is a structure in the form* $p(t_1, \ldots, t_n)$*, with* $p \in PRED$ *and* $t_1, \ldots, t_n \in TERMS$*.*

Predicates are also written in lower-case (e.g. `student(john)`).

Definition 4 (Literal). *The set of* literals *L consists of* positive literals *(atoms),* negative literals *(negated atoms):*

- $L = L^+ \cup L^-$
- $L^+ = \{p(t_1, \ldots, t_n) \mid p \in PRED,\ t_i \in TERM,\ n > 0\}$
- $L^- = \{-p(t_1, \ldots, t_n) \mid p \in PRED,\ t_i \in TERM,\ n > 0\}$

where '−' stands for strong negation.[11]

Definition 5 (Extended Literal). *The set of* extended literals L^* *consists of literals and* default negation literals*:*

- $L^* = L \cup L^{\text{not}}$
- $L^{\text{not}} = \{\text{not}\, l \mid l \in L\}$

*where '*not*' stands for* default negation.[12]

Definition 6 (Ground Set). *For all set of extended literals* $X \subseteq L^*$ *we denote with* X_{grnd} *the* ground set *of X, i.e. the elements of X which do not contain any variable.*

Definition 7 (Variable Set). *We denote with* $varOf : L^* \to VAR$*, the function returning the variables used in a extended literal.*

[11] Strong negation is used to reify an explicitly false situation (e.g. "It does not rain").
[12] Default negation is used to reify a situation in which something cannot be retrieved/inferred (e.g. 'It is unknown whether it rains or not').

3.2 Net topology

Definition 8 (Net). *A* net *is a* bipartite directed graph *connecting two finite sets of nodes, called* places *and* transitions. *It can be written as* $N = \langle P, T, E \rangle$, *where:*

- $P = \{p_1, \ldots, p_n\}$ *is the set of* place *nodes;*
- $T = \{t_1, \ldots, t_m\}$ *is the set of* transition *nodes;*
- $E = E^+ \cup E^-$ *is the set of* arcs *connecting them:* $E^+ \subseteq T \times P$ *are the arcs going from transitions to places,* $E^- \subseteq P \times T$ *are the arcs going from places to transitions.*

We will denote an arc between two nodes n_i *and* n_j *as* (n_i, n_j).

Definition 9 (Declarative Net of Places). *A* declarative net of places *is a* bipartite directed graph *connecting two finite sets of nodes:* places *labeled with literals and* logic operator nodes. *It can be written as* $LPN = \langle P, C_P, LP, C_{LP}, DE_{LP} \rangle$, *where:*

- $P = \{p_1, \ldots, p_n\}$ *is the set of* place *nodes;*
- $C_P : P \to L^*$ *is a labeling functions, associating literals to places;*
- $LP = \{l_1, \ldots, l_m\}$ *is the set of* logic operator *nodes applied on places;*
- *Given the set of logic operators for places* $OP_{LP} = \{\neg, -, \wedge, \vee, \rightarrow, \leftrightarrow, \ldots\}$, $C_{LP} : LP \to OP_{LP}$ *maps logic operators to logic operator nodes.*
- $DE_{LP} = DE^+_{LP} \cup DE^-_{LP}$ *is the set of arcs connecting logic operator nodes to places:* $DE^+_{LP} \subseteq LP \times P$ *from logic operator nodes to places,* $DE^-_{LP} \subseteq P \times LP$ *from places to logic operator nodes;*

Definition 10 (Declarative Net of Transitions). *A* declarative net of transitions *can be written as* $LTN = \langle P, C_P, T, C_T, LT, C_{LT}, DE_{LT} \rangle$, *where:*

- $P = \{p_1, \ldots, p_n\}$ *is the set of* place *nodes;*
- $C_P : P \to L^*$ *is a labeling function, associating literals to places;*
- $T = \{t_1, \ldots, t_n\}$ *is the set of* transition *nodes;*
- $C_T : T \to L$ *is a labeling function, associating literals to transitions;*
- $LT = \{l_1, \ldots, l_m\}$ *is the set of* logic operator *nodes applied on transitions;*
- *Given the set of logic operators for transitions* $OP_{LT} = \{ \rightarrow, \leftrightarrow \}$, $C_{LT} : LT \to OP_{LT}$ *maps logic operators for transitions to logic operator nodes.*
- $DE_{LT} = DE^+_{LT} \cup DE^-_{LT}$ *is the set of arcs connecting logic operator nodes to transitions:* $DE^+_{LT} \subseteq LT \times T$ *from logic operator nodes to transitions,* $DE^-_{LT} \subseteq (T \cup P) \times LT$ *from transitions* and places *to logic operator nodes;*

The declarative nets of places and of transitions operate at different ontological levels. The first is used to compute the *steady state* of the system. The second connects all *impulse events* that are in a relation of dependency. As, in our case, the Petri Net is associated with an interleaving semantics, the execution cycle cannot autonomously originate multiple events; therefore, the only operators available at event level are implication and bi-implication.

Definition 11 (Logic Programming Petri Net). *A* Logic Programming Petri Net *LPPN is a Petri Net whose places and transitions are labeled with literals, enriched with declarative nets of places and of transitions. It can be written as* $LPPN = \langle P, C_P, T, C_T, PE, LP, C_{LP}, DE_{LP}, LT, C_{LT}, DE_{LT}, M \rangle$*, where:*

- $\langle P, T, PE \rangle$ *is a net; PE stands for* procedural edges*;*
- $C_P : P \to L^*$ *and* $C_T : T \to L$ *are the labeling functions, associating literals to places and to transitions, respectively;*
- $\langle P, C_P, LP, C_{LP}, DE_{LP} \rangle$ *is a declarative nets of places;*
- $\langle P, C_P, T, C_T, LT, C_{LT}, DE_{LT} \rangle$ *is a declarative nets of transitions;*
- $M : P \to 2^{L^*_{grnd}}$ *returns the marking of a place, i.e. the set of labels of the tokens that the place contains.*

From this definition, we observe the following properties:

- if $LP \cup LT = \varnothing$, we have a *strictly procedural* LPPN;
- if $T = \varnothing$, we have a *strictly declarative* LPPN.

A strictly declarative LPPN can be directly mapped to an ASP program.

3.3 Syntaxic constraints on topology

We derive from ASP the following constraints:

- C_P cannot map output places to default negation literals; in ASP, negation-as-failure literals can be used only in the body of rules: we cannot conclude *ignorance* (default negation literals) but, if necessary for the modeler, only *undecidability* in terms of positive or negative polarity (null literals).
- C_T cannot map transitions to default negation literals (L instead of L^*);
- if $n \in LP$ is mapped to:
  - a negation ($\neg$, $-$), it requires only one input place;
  - a conjunction or a disjunction ($\wedge$, $\vee$), it requires at least one input place;
  - a negation, a conjunction or a disjunction ($\neg$, $-$, $\wedge$, $\vee$) it requires at least one output place;
- if $n \in LP \cup LT$ is mapped to:
  - an implication ($\rightarrow$), it requires at least one input *or* one output node;
  - a bi-implication ($\leftrightarrow$), it requires one input *and* one output node;

Furthermore, to ensure that the execution is valid, there should be always a connection of parameters between outputs and inputs, in order to transport only ground elements: More formally:

Definition 12 (Preset, Postset). *Given a node* x*, the set of nodes connected as* inputs *to* x *is called the* preset *of* x*; in formula,* $\bullet x = \{y | (y, x) \in E\}$*. Similarly, the set of nodes connected as* outputs *to* x *is called the* postset *of* x*; in formula,* $x\bullet = \{y | (x, y) \in E\}$*.*

Definition 13 (Safeness conditions). *For each transition, all variables that are present in the labels of output places have to be present in one of the input places, or in the transition label:*

$$\begin{gathered}\forall t \in T, \forall p_o \in t\bullet, \\ \forall v \in varOf(C_P(p_o)) : v \in varOf(C_T(t)) \vee \exists p_i \in \bullet t : v \in varOf(C_P(p_i))\end{gathered} \tag{1}$$

For each logic operator node, all variables that are present in the labels of output nodes have to be present in at least one of the input nodes:

$$\forall n \in LP, \forall p_o \in l\bullet, \forall v \in varOf(C_P(p_o)), \exists p_i \in \bullet n : v \in varOf(C_P(p_i)) \tag{2}$$

$$\begin{gathered}\forall n \in LT, \forall t_o \in l\bullet, \forall v \in varOf(C_T(t_o)), \\ \exists t_i \in \bullet n \cap T : v \in varOf(C_T(t_i)) \vee \exists p_i \in \bullet n \cap P : v \in varOf(C_P(p_i))\end{gathered} \tag{3}$$

3.4 Hybrid operational semantics

As LPPN transitions are labeled, their firing forges a *transition event*. However, we need to distinguish the *direct*, external firings (started by the simulator/interpreter) from the *indirect*, immediately propagated, internal firings (triggered by the declarative net of transitions).

Execution cycle The execution cycle of a LPPN consists of three-steps. Given a certain marking M:

1. an enabled transition is selected to *pre-fire*, so determining a 'source' *transition event*;
2. all propagations of this event are entailed using the specifications of the declarative net of transitions, obtaining a set of *transition events*;
3. all transition events are fired, producing and consuming the relative tokens.[13]

Running example Let us consider the LPPN in Fig. 4. The net provides causal mechanisms, declarative constraints, and only one entity—`c1(a)`—localized in p_1. The following execution paths are possible:

- t_2 fires the transition event `e2`, consuming `c1(a)` from p_1;
- t_1 fires the transition event `e1(a)`, consuming `c1(a)` from p_1 and producing `c2(a)` in p_2; thanks to dt_1, the existence of `c2(a)` is a sufficient reason for immediately reifying `c5(a)`.

Restarting again, let us suppose the presence of an additional element `c4(b)` in p_4, enabling t_3. With this variation, the second execution path would be slightly different because:

- thanks to dp_1, the transition event `e1(a)` propagates to t_3, that in turn, being enabled, immediately fires another event `e3(a)`, consuming `c4(b)` and producing `c3(a, b)` in p_3.

[13] This proposal considers an interleaved semantics for the selection of the transition (pre-firing), and a reactive, deterministic firing. Future extensions to non-deterministic cases are however possible.

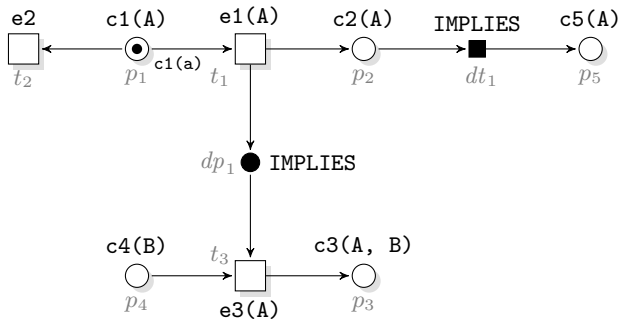

Fig. 4: An example of LPPN with all the elements provided by the notation.

Formalization A crucial element of the execution mechanism is the selection of an adequate substitution of variables, built upon the locally available (information of) conditions and events.

Definition 14 (Variable substitution). *A substitution B is a set of associations between variables and constants: $B \subseteq VAR \times CONST$. Given a literal l, we denote with $l_{[B]}$ the rewriting of l, substituting all variables v in l with the constant associated with v in B. Similarly, given a set X of literals, $X_{[B]}$ is the set obtained by substituting with B all literals of X.*

Definition 15 (Enabled transition). *A transition t is* enabled *in a marking M if it is possible to construct a variable substitution using the tokens available in the input places:*

$$\begin{aligned} Enabled(t) \equiv \exists B \subseteq VAR \times CONST : \\ \forall p_i \in \bullet t, C_P(p)_{[B]} \in M(p) \end{aligned} \tag{4}$$

We define with $Subs(t) = \{B_1, \ldots\}$ the set of possible substitutions (depending on the marking as well). Equivalently, if $Subs(t) = \varnothing$ the transition is not enabled.

In LPPNs, transitions have labels has well. When they are fired they forge *transition events*, volatile elements whose label specifies the actual event. Transition events are therefore always associated with a certain source transition.

Definition 16 (Event label). *Given a transition event e, we denote with $C_E(e) \in L_{grnd}$ the label describing e.*

For safeness reasons, (direct) firing has to be disabled if the transition label would remain with some free variable after substitution. We define therefore a dedicated additional enabling condition, that can be acknowledged off-line, via syntaxic analysis.

Definition 17 (Safe transition). *A transition t is* safe *if all variables of its literal label are included in the literal label of its input places:*

$$Safe(t) \equiv \forall v \in varOf(C_T(t)), \exists p_i \in \bullet t : v \in varOf(C_P(p_i)) \tag{5}$$

Given a set of transitions T, $Safe(T)$ is the set of elements of T that are safe.

Safe transitions may pre-fire, producing a *transition event* with a substitution selected amongst those available.

Definition 18 (Pre-firing). *A safe enabled transition t* pre-fires *in a marking M by selecting one substitution B from $Subs(t)$, and* forging *with B a* transition event*:*

$$\forall t \in Enabled(T) \cup Safe(T) :$$
$$t \text{ pre-fires} \equiv \exists B \in Subs(t), \exists e \in TransitionEvent(t) : C_E(e) = C_T(t)_{[B]}$$

Where $TransitionEvent(t)$ is the set of transition events fired by transition t.

At this point, the *transition event* forged with the pre-firing is propagated through the declarative net of transitions.

Definition 19 (Firing). *An enabled transition t* fires *by propagation, selecting (if it exists) one substitution B from $Subs(t)$ compatible with the incoming transition event e,* consuming *a token satisfying B from each input place, and* forging *with B a token in each output place:*

$$\forall t \in Enabled(T) :$$
$$t \text{ fires} \equiv \exists e \in TransitionEvent(t), \exists B \in Subs(t) :$$
$$C_T(t)_{[B]} = C_E(e) \wedge$$
$$\forall p_i \in {\bullet t} : M'(p_i) = M(p_i) \setminus C_P(p_i)_{[B]} \wedge$$
$$\forall p_o \in t{\bullet} : M'(p_o) = M(p_o) \cup C_P(p_o)_{[B]}$$

3.5 Denotational semantics

One of the possibilities to validate a formal language is to map it into another formal language, i.e. to provide a *denotational semantics*. A LPPN consists of two components. The declarative component, by construction, can be directly rewritten as ASP code. As we are already halfway down the path, we can translate the remaining procedural component into ASP code as well.

Event Calculus axioms An existing solution to treat change in logic programming is *event calculus* (EC) [17,30]. The simple version is already satisfactory for our purposes. A modification of the original axioms is, however, necessary to deal with the *locality* brought by places and transitions:

```
holdsAt(F, P, N) :-
   initially(F, P), not clipped(0, F, P, N),
   fluent(F), place(P), time(N).

holdsAt(F, P, N2) :-
   firesAt(E, T, N1), N1 < N2,
```

```
   initiates(T, F, P, N1), not clipped(N1, F, P, N2),
   place(P), transition(T), fluent(F),
   event(E), time(N1), time(N2).

clipped(N1, F, P, N2) :-
   firesAt(E, T, N), N1 <= N, N < N2, terminates(T, F, P, N),
   place(P), transition(T), fluent(F),
   event(E), time(N1), time(N2), time(N).
```

Operational semantics axioms Considering interleaved semantics for the selection of the transition, and for simplicity, only deterministic firing, we have the following rules:

i. all safe enabled transitions may or may not pre-fire,
ii. at least one transition must pre-fire per step, i.e. it is impossible that no transition fire, save the last step,
iii. at maximum one transition can pre-fire per step.

In ASP code:

```
{prefiresAt(T, N)} :-                                    % (i)
   enabled(T, N), safe(T), transition(T), time(N).

someTransitionPrefiresAt(N) :-                           % (ii)
   prefiresAt(T, N), transition(T), time(N).
:- not someTransitionPrefiresAt(0).
:- N > 0, not someTransitionPrefiresAt(N-1), time(N).

:- prefiresAt(T1, N), prefiresAt(T2, N), T1 != T2,       % (iii)
   transition(T1), transition(T2), time(N).
```

Transformation of a LPPN to an ASP program The mapping of a given LPPN to an equivalent ASP program includes the previous axioms and the output of the following steps:

i. for each place p, with label $C_P(p)$
   (a) type it as place,
   (b) define its label as fluent,
   (c) specify its initial state,
   (d) for each place with more than one output, write down that you cannot consume more than the only available token.
ii. for each transition t, with label $C_T(t)$
   (a) type it as transition,
   (b) define its label as event,
   (c) specify if it is safe,
   (d) define the conditions for which it is enabled,

(e) if it is safe, define the substitution transforming the prefiring in firing,
(f) define how the variable substitution is defined
(g) for each output place, define how to create tokens in the output places,
(h) for each input place, define how to consume tokens in the output places.

iii. for each logic operator node *dt* for places,
(a) specify the logic constraint to be applied between inputs and outputs.

iv. for each logic operator node *dp* for transitions,
(a) write down the logic dependencies between transitions allowing the propagation.

As a concrete example, we apply these actions on some of the components of the LPPN in Fig. 4:

```
%%% p1
place(p1).                                                          % 1.a
fluent(c1(A)) :- a(A).                                              % 1.b
initially(c1(a), p1).                                               % 1.c
:- 2{terminates(t1, c1(A), p1, N); terminates(t2, c1(A), p1, N)},  % 1.d
   a(A), time(N).

%%% t1
trans(t1).                                                          % 2.a
event(e1(A)) :- a(A).                                               % 2.b
safe(t1).                                                           % 2.c
enabled(t1, N) :- holdsAt(c1(A), p1, N), a(A), time(N).             % 2.d
firesAt(e1(A), t1, N) :-                                            % 2.e
   holdsAt(c1(A), p1, N), prefiresAt(t1, N), a(A), time(N).
appliesAt(t1, subs(A), N) :-                                        % 2.f
   holdsAt(c1(A), p1, N), firesAt(e1(A), t1, N), a(A), time(N).
initiates(t1, c2(A), p2, N) :-                                      % 2.g
   appliesAt(t1, subs(A), N), a(A), time(N).
terminates(t1, c1(A), p1, N) :-                                     % 2.h
   appliesAt(t1, subs(A), N), a(A), time(N).

%% dt1
holdsAt(c5(A), p5, N) :- holdsAt(c2(A), p2, N), a(A), time(N).      % 3.a

%% dp1
firesAt(e3(A), t3, N) :-                                            % 4.a
   holdsAt(c4(B), p4, N), appliesAt(t1, e1(A), N), a(A), time(N).
```

Output With the transformation steps given above, valid LPPNs can be transformed into ASP programs, whose answer sets, defining a temporal range (with the instruction "`time(0..`n`).`") represent all possible executions path after at most n steps. For instance, running our example[14] with the ASP engine `clingo` [8], we obtain:

[14] The full code of the example is available online at `http://justinian.leibnizcenter.org/lppn2clingo/`; it can be run within the browser at: `https://potassco.org/clingo/run/`.

```
clingo version 5.2.0
Reading from stdin
Solving...
Answer: 1
holdsAt(c1(a),p1,0) holdsAt(c4(b),p4,0) holdsAt(c4(b),p4,1)
firesAt(e2,t2,0)
Answer: 2
holdsAt(c1(a),p1,0) holdsAt(c4(b),p4,0) firesAt(e1(a),t1,0)
firesAt(e3(a),t3,0) holdsAt(c2(a),p2,1) holdsAt(c3(a,b),p3,1)
holdsAt(c5(a),p5,1)
SATISFIABLE
```

4 Evaluation

The proposal presented above is currently used for developing a prototype application for specifying, executing and analyzing LPPNs[15]. A previous prototype was being developed in Java, but the repeated external calls to any ASP solver (typically developed in C++) were prohibitive in computational terms. Fortunately, `clingo` has introduced quite recently an API for Python, enabling a direct control of the solver instance, without regrounding the program at each cycle. Interestingly, this enabled us to perform some direct evaluation: the resulting modular architecture reproduces the partial overlap of the computational efforts in the two operationalizations. When we refer to the denotational semantics, the solver intervenes *fully* to solve the event calculus program. Instead, when we refer to the hybrid operational semantics, the solver intervenes *partially*, twice, to solve the declarative components of the net, while the rest of the computational burden is on the module responsible for the Petri Net execution. In this context, one might ask if we can observe some patterns in the performances between the alternative modes of analysis/execution.

At the moment, we have only evaluated a propositional version of LPPN, and a limited series of structures, namely compositions of minimal *serial* elements (a transition with an input and output places) or minimal *forking* elements (a place with two output transitions). In order to implement the procedural component of the operational semantics, the current Petri Net analysis module builds upon a *brute force* execution algorithm, and *depth-first* search with *backtracking* to cover all the possible execution paths.

Table 1 summarizes the performances of 10 executions of different configurations of nets.[16] The data partially confirms what we expected: the analysis based on the operational semantics (BE+BT) clearly outperforms and scales excellently for the serial configurations, while that based on the denotational semantics (EC) scales poorly for such a simple structure. However, for the nested forking configurations, BE+BT is evidently slower; intuitively this is due to the

[15] Available at `http://github.com/s1l3n0/pypneu`.

[16] The tests were run on a laptop provided with a 1.5 GHz Quadcore processor and 3Gb RAM.

	depth of composition of minimal structures							
serial	1	6	11	16	21	26	31	36
EC	0.0003	0.0023	0.0106	0.0324	0.0805	0.1663	0.2994	0.4961
BE+BT	0.0009	0.0032	0.0067	0.0121	0.0187	0.0256	0.0333	0.0437
forking	1	2	3	4	5	6	7	8
EC	0.0008	0.0016	0.0043	0.0114	0.0362	0.1033	0.2964	1.0137
BE+BT	0.0013	0.0033	0.0096	0.0318	0.1126	0.4396	1.7806	7.9581

Table 1: Average execution time (ms) for the analysis of different configurations of propositional LPPNs, performed following the hybrid operational semantics, via *brute execution* and *backtracking* (BE+BT); or the denotational semantics, via *event calculus* (EC).

efficient search and pruning capabilities of ASP. However, if we consider the incremental ratio, we observe that it is roughly comparable to the results provided by the second mode. In effect, the Python code of the executor/analyzer is not optimized as `clingo`, on the contrary, for many aspects this represents a lower-bound on the possible implementation choices.

Conclusion and further developments

The paper presents an extension to Petri Net that was elaborated with a practical goal in mind: a modeling notation relatively simple to be presented to non-experts, that could handle explicit knowledge (terminological, or more in general, declarative), and that could model causation/temporality. Further discussion on these aspects and associated examples of modeling (of norms, of narratives, of agent reasoning, etc.) using LPPN can be found in [31]. Here, the focus is on the general characteristics of the notation, on its consolidation, and on the associated computational aspects. Many aspects remain to be investigated, though. The propagated firing (which is plausibly a unique characteristic of LPPN with respect to other extensions) is evidently an exception to the *interleaving semantics* of standard Petri Nets. However, the declarative dimension for events allows us to treat at higher abstraction all phenomena for which there is a viable specification at outcome level. Consider terminological or institutional qualifications issued from a narrative model: for the model to work, it is not important how those epistemic inferences are performed, but only that their result is logically consistent. This is in analogy with how engineers distinguish the study on steady states, from the study on transients. However, from a theoretical point of view, a generic LPPN may still present configurations incompatible with the duality of procedural and declarative dimensions. For instance, when propagated firings involve transitions sharing the same input places, they may consume more than what is available in the input tokens, going against the rules of the execution

semantics. If this occurs, this is plausibly a consequence of an error in the modeling phase. For the moment, we simply raise a warning to the modeler. A more structured solution remains to be investigated.

Future developments concern the extension and the refinement of the evaluation, considering the full predicate notation and mixed networks (of declarative, procedural components) with mixed configurations (serial, forks, joins, etc.). The actual impact on real models should be evaluated as well: scenarios describing cases have very few forks, they rather function as *orchestrated* (i.e. directed from the scenario) *scripts* (distributed amongst actors). Consequently, applications that require the use of scenarios (e.g. for interpretation, model-based diagnosis, conformance checking, etc.) may take advantage of the hybrid operational semantics. The computational improvement may be further extended considering existing proposals in the literature. For instance, execution algorithms alternative to brute execution [25,27]; or decomposition techniques, for instance in *single-entry-single-exit* (SESE) components [26], that open up the possibility of concurrent execution. Finally, these results should be confronted with existing techniques for handling temporal reasoning and causality, e.g. the already cited *Action languages* [11], related works (e.g. F2LP [19]) and applications (CCalc, Coala, Cplus2ASP); optimized versions of Event Calculus (e.g. [4]); applications based on LTL, CTL and related formalisms.

References

1. Anderson, J.R.: Learning and memory: an integrated approach. John Wiley & Sons, Inc. (1995)
2. Anwar, S., Baral, C., Inoue, K.: Encoding Higher Level Extensions of Petri Nets in Answer Set Programming. Logic Programming and Nonmonotonic Reasoning - Lecture Notes in Computer Science (2013)
3. Anwar, S., Baral, C., Inoue, K.: Encoding Petri nets in answer set programming for simulation based reasoning. TPLP 13(4-5-Online-Supplement) (2013)
4. Artikis, A., Sergot, M., Paliouras, G.: An Event Calculus for Event Recognition. IEEE Transactions on Knowledge and Data Engineering 27(4), 895–908 (2015)
5. Behrens, T.M., Dix, J.: Model checking multi-agent systems with logic based Petri nets. Annals of Mathematics and Artificial Intelligence 51(2-4), 81–121 (feb 2008)
6. Bonner, A.J., Kifer, M.: Transaction Logic Programming. International Conference on Logic Programming pp. 257–279 (1993)
7. Breuker, J., Hoekstra, R.: Core concepts of law: taking common-sense seriously. Proc. of Formal Ontologies in Information (2004)
8. Eiter, T., Faber, W., Fink, M., Woltran, S.: A user's guide to gringo, clasp, clingo, and iclingo. Annals of Mathematics and Artificial Intelligence 51(2-4), 123–165 (2008)
9. Ferraris, P., Lee, J.: Representing first-order causal theories by logic programs. Theory and Practice of Logic Programming 12(03), 383–412 (may 2012)
10. Gebser, M., Grote, T., Schaub, T.: Coala: A compiler from action languages to ASP. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 6341 LNAI, 360–364 (2010)

11. Gelfond, M., Lifschitz, V.: Action languages. Electronic Transactions on AI (1998)
12. Gelfond, M., Lifschitz, V.: The stable model semantics for logic programming. Proceedings of International Logic Programming Conference and Symposium pp. 1070–1080 (1988)
13. Genrich, H.J.: Predicate/Transition Nets. In: Proceedings Advances in Petri nets 1986. pp. 207–247 (1987)
14. Jensen, K.: Coloured Petri Nets: Basic Concepts, Analysis Methods and Practical Use. Springer-Verlag, London, UK (1996)
15. Kemmerer, D., Eggleston, A.: Nouns and verbs in the brain: Implications of linguistic typology for cognitive neuroscience. Lingua 120(12), 2686–2690 (2010)
16. Kowalski, R., Sadri, F.: Integrating logic programming and production systems in abductive logic programming agents. Web Reasoning and Rule Systems LNCS 5837, 1–23 (2009)
17. Kowalski, R., Sergot, M.: A logic based calculus of events. New Generation Computing 4(June 1975), 67–95 (1986)
18. van Lambalgen, M., Hamm, F.: The Proper Treatment of Events. Wiley-Blackwell, Oxford, UK (2004)
19. Lee, J., Palla, R.: System f2lp - computing answer sets of first-order formulas. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 5753 LNAI, 515–521 (2009)
20. Lefevre, C., Nicolas, P.: A First Order Forward Chaining for Answer Set Computing. LPNMR 2009 LNCS 5753, 196–208 (2009)
21. de Leoni, M., Munoz-Gama, J., Carmona, J., van der Aalst, W.M.P.: Decomposing Alignment-Based Conformance Checking of Data-Aware Process Models. In: Meersman, R., Panetto, H., Dillon, T., Missikoff, M., Liu, L., Pastor, O., Cuzzocrea, A., Sellis, T. (eds.) On the Move to Meaningful Internet Systems: OTM 2014, pp. 3–20. Springer Berlin Heidelberg, Berlin, Heidelberg (2014)
22. Lifschitz, V.: Two-Valued Logic Programs (2012)
23. McCarthy, J., Hayes, P.J.: Some philosophical problems from the standpoint of artificial intelligence. In: Machine Intelligence, pp. 1–51. Edimburgh University Press (1969)
24. Moore, R.C.: Semantical Considerations on Nonmonotonic Logic. Artificial Intelligence 25(1), 75–94 (jan 1985)
25. Moreno, R.P., Salcedo, J.L.V.: Performance evaluation of petri nets execution algorithms. Conference Proceedings - IEEE International Conference on Systems, Man and Cybernetics pp. 1400–1407 (2007)
26. Munoz-Gama, J., Carmona, J., Van Der Aalst, W.M.P.: Single-Entry Single-Exit decomposed conformance checking. Information Systems 46, 102–122 (2014)
27. Piedrafita, R., Villarroel, J.L.: Performance evaluation of petri nets centralized implementation. The execution time controller. Discrete Event Dynamic Systems: Theory and Applications 21(2), 139–169 (2011)
28. Reiter, R.: A logic for default reasoning. Artificial Intelligence 13(1-2), 81–132 (apr 1980)
29. Reiter, R.: Knowledge in action: logical foundations for specifying and implementing dynamical systems (2001)
30. Shanahan, M.: The event calculus explained. Artificial intelligence today pp. 409–430 (1999)
31. Sileno, G.: Aligning Law and Action. Ph.D. thesis, University of Amsterdam (2016)
32. Thielscher, M.: From situation calculus to fluent calculus: State update axioms as a solution to the inferential frame problem. Artificial Intelligence 111(1-2), 277–299 (jul 1999)